

"Express Mail" mailing label number:

EV 335896263 US

**RING OVERLAY NETWORK DEDICATED TO CARRY BROADCAST
TRAFFIC TO DSLAMS**

Albert Chungbor Wan
Ahmad Ansari
David Hartman
Mehrad Yasrebi

Field of the Disclosure

[1001] The present disclosure relates to multicast telecommunications.

Description of the Related Art

[1002] FIG. 1 is a simplified view of an Internet Protocol (IP) network 10. The IP network 10 provides access to the Internet 12 for residential users 14 and corporate users 16. IP networks 10 and the Internet 12 were designed such that various network elements such as routers 20, Asynchronous Transfer Mode (ATM) routers 22, switches 24, access devices such as Customer Premises Equipment (CPEs) 26 and 30, and hosts logically communicate using transport data packets in a pair-wise or unicast fashion.

[1003] Many Internet access services are offered by telephone companies (telcos) using technologies such as various types of Digital Subscriber Lines (DSL), which are generalized as xDSL. An example of an xDSL is Asymmetric DSL (ADSL). In xDSL, carrier-side access devices are called DSL Access Multiplexers (DSLAMs). DSLAMs perform traffic aggregation, grooming, policing, access multiplexing and other tasks.

[1004] FIG. 2 shows placement of DSLAMs in relation to the IP network 10 and the Internet 12. Each of DSLAMs 32, 34, 36 and 40 is linked to one of the router 20 or 22. The DSLAM 32 is disposed at a central office 42. The DSLAMs 34, 36 and 40 are disposed at a central office 44. The DSLAM 32 provides ADSL access to multiple customer premises 46. The DSLAM 34 provides ADSL access to multiple customer

premises 50. The DSLAM 36 provides ADSL access to multiple customer premises 52. The DSLAM 40 provides ADSL access to multiple customer premises 54.

[1005] CPEs at the customer premises have DSL modems to interface with the DSLAMs to provide access to the Internet 12. Users can perform various Internet applications such as browsing, file transfer, peer-to-peer communication and instant messaging. The aforementioned applications, which involve sending data from a sender to a recipient in a one-to-one unicast relationship, are efficiently performed by the Internet.

[1006] The maturity of distributed computing necessitates the need to accommodate non-unicast communication models. An example is broadcast video applications, i.e. applications that are similar in function to terrestrial broadcast video, cable video and satellite video. Such applications require distribution of data from a single source to multiple recipients at multiple destinations. A different approach to network design is needed to efficiently handle multicast applications.

[1007] The Internet Engineering Task Force (IETF), which is a loosely-formed Internet-standard-setting organization, has several IP multicast routing protocols that can be used by service providers to emulate broadcast on a unicast-based network. Examples of the protocols include Protocol-Independent Multicast (PIM) sparse mode, PIM dense mode and Multicast Open Shortest Path First (MOSPF). Using these multicast routing protocols, routers and other network elements that understand routing protocols in the unicast network can exchange multicast routing information and form IP multicast distribution trees. However, only the routers that participate in the multicast distribution tree provide the IP multicast data transport.

[1008] Depending on the routing protocol that is used, the creation of the multicast distribution tree can be either dynamic or static. The tree-building process comprises routers listening for IGMP join and leave requests from the CPEs, and performing associated tree pruning and grafting processes between routers. IGMP, which stands for Internet Group Management Protocol, is a mechanism used by IP-multicast-capable CPE devices to signal their channel group join/leave interest to a nearest network element.

The network element is usually a routing device configured to understand multicast routing protocols and to participate in the multicast routing decisions as described above.

[1009] None of these IP multicast routing protocols provide intrinsic end-to-end Quality of Service (QoS) guarantees. However, a limited IP Class of Service (CoS) mechanism does exist for the unicast model. The lack of end-to-end QoS capability from the IP layer leaves one to use lower layer protocols (e.g. ATM at layer 2), or to unnecessarily over-engineer the IP network for abundant/underutilized bandwidth between routers and switches to guarantee end-to-end QoS.

[1010] Over-engineering requires both links and routers (or switches) in the network to have excess bandwidth and routing capacity to handle peak traffic usage scenarios. To address the inherent inefficiency associated with over-engineering, statistical multiplexing models are used to share the over-engineered transport capacity among video channels.

[1011] While most unicast-based Internet applications are Unspecified Bit Rate (UBR)-based, broadcast entertainment video applications are mostly Constant Bit Rate (CBR)-based or Variable Bit Rate (VBR)-based. Many audio/video compression technologies, including MPEG2, can have a peak data rate much higher than the normal data rate. A high peak data rate occurs, for example, in a transition from a talking head with a simple background to a high-motion scene with a detailed background. Over-engineering without taking advantage of statistical multiplexing can quickly erode a price advantage that IP infrastructure brings to service providers.

[1012] Some small telcos are offering, for exploratory trial purposes, broadcast video services on their xDSL access network by expanding their unicast-centric network to support multicast. Some consider these solutions as being complicated, inflexible, difficult to expand, and cost ineffective. In these approaches, video content from satellite feeds and content servers in a head-end office is transported through switched legacy data networks to reach all video viewers. Because small telcos usually have smaller networks and smaller subscriber populations, their network normally has fewer element hop counts from the head-end to the CPEs and is easier to over-engineer. Over-engineering of larger

telcos' networks, however, would not result in a feasible, scalable, and cost-effective solution.

BRIEF DESCRIPTION OF THE DRAWINGS

[1013] The present invention is pointed out with particularity in the appended claims. However, other features are described in the following detailed description in conjunction with the accompanying drawing in which:

[1014] FIG. 1 is a simplified view of an Internet Protocol network;

[1015] FIG. 2 shows placement of DSLAMs in relation to the IP network and the Internet;

[1016] FIG. 3 is a schematic, block diagram of an embodiment of an architecture in accordance with the present invention;

[1017] FIG. 4 is a schematic, block diagram of an enhancement to the embodiment of FIG. 3;

[1018] FIG. 5 is a flow chart of an embodiment of a method of providing broadcast video content using embodiments of the architecture; and

[1019] FIG. 6 is a flow chart of an embodiment of a method of transporting unicast interactions using embodiments of the architecture.

DETAILED DESCRIPTION OF THE DRAWINGS

[1020] Embodiments of the present invention provide network architectural changes and an associated control mechanism that allows Internet access providers to build a transport network for entertainment broadcast video with improved QoS and cost savings over the traditional IETF IP multicast proposals. The architecture includes an overlay broadcast network that co-exists with an existing legacy data network. The overlay broadcast

network can use SONET technology or an alternative technology to unidirectionally transport broadcast video content. The associated control mechanism for transporting the broadcast video content is used in addition to existing mechanisms for the legacy data network.

[1021] FIG. 3 is a schematic, block diagram of an embodiment of an architecture in accordance with the present invention. A legacy data network 100 communicates interactive data exchanges between a head-end network 102 and multiple CPEs 104 and 106. Although two CPEs are illustrated, those having ordinary skill will appreciate that more than two CPEs can be served using the architecture.

[1022] The head-end network 102 comprises a satellite dish/receiver 110 and a video-on-demand (VOD) server farm 112 to provide video programming. The satellite dish/receiver 110 and the VOD server farm 112 are linked to a router 114 to provide network access to the video programming. The head-end network 102 may further comprise a computer 116 and a database 120 to provide back office and access control features.

[1023] The legacy data network 100 comprises an IP/ATM network 122 in communication with the head-end network 102 and the Internet 124. Carrier-side access to the IP/ATM network 122 is provided by a DSLAM 126. The legacy data network 100 comprises an existing unicast network, such as a legacy ADSL data network.

[1024] An overlay network 130 augments the legacy data network 100. Preferably, the overlay network 130 comprises a SONET network having at least one SONET ring. A multi-tiered SONET ring architecture can be used to enhance the coverage area for the overlay network, and/or to provide intra-LATA/inter-LATA coverage. For purposes of illustration and example, the overlay network 130 is depicted to comprise a first-tier SONET ring 132 and a tier-n SONET ring 134. The SONET rings 132 and 134 are connected by a cross connect element 136.

[1025] The SONET network is provisioned to carry only broadcast traffic from the head-end network 102 to the access networks. In contrast, the interactive data exchanges

between the head-end network 102 and the access networks, which are easily accommodated using the unicast model, are carried by the legacy data network 100. The SONET network has improved performance of providing broadcast video from the head-end network 102 to the DSLAM 126.

[1026] FIG. 4 is a schematic, block diagram of an enhancement to the embodiment of FIG. 3. In this embodiment, a dedicated separate data network 140 uses a unicast model to carry interactive data and control changes between the head-end network 102 and the access networks. Using the dedicated separate data network 140 instead of the legacy data network 100 improves the performance of interactive data and control exchanges. Examples of data that would be carried on the dedicated separate data network 140 include quality-sensitive applications such as exchanges to support interactive television and VOD.

[1027] In one embodiment, the dedicated separate data network 140 comprises a Virtual Private Network (VPN). The deployment of the VPN can be made using any VPN technology that provides the performance and QoS needed by a service provider. The VPN can be provided by an IP/ATM network 142.

[1028] FIG. 5 is a flow chart of an embodiment of a method of providing broadcast video content in accordance with the present invention. The method applies to the architectures of both FIG. 3 and FIG. 4 unless otherwise indicated.

[1029] As indicated by block 200, the method comprises provisioning at least one SONET ring, such as the SONET rings 132 and 134, with Virtual Circuits (VCs) from an ingress Add-Drop Multiplexer (ADM) 202 to each of multiple egress ADMs 204. This act can be performed using a suitable scheme such as Unidirectional Path-Switched Ring (UPSR). UPSR is an example of suitable use of network bandwidth in transporting unidirectional broadcast video to DSLAMs. Alternatively, a different scheme may be used as selected by the service provider.

[1030] The ingress ADM 202 is used to receive multiple broadcast video contents from the head-end network 102. The egress ADMs 204 are used to feed the broadcast video

contents to DSLAMs, including the DSLAM 126, in the access network. An egress ADM 208 connects to a trunking interface (i.e. a network interface) of the DSLAM 126. Only one egress ADM is needed for a group of DSLAMs. If more than one DSLAM is in the group, the broadcast video data can be replicated using inexpensive IP layer-2 switches between the egress ADM and the receiving DSLAMs. Although optional, the aforementioned use of the layer-2 switch provides a cost saving approach for multi-DSLAM scenarios. If preferred, the service provider can use one egress ADM to one DSLAM.

[1031] As indicated by block 210, a mechanism is provided in the DSLAM 126 or alternative network element to bypass a traditional multicast routing mechanism and to provide a different way of multicast routing through the overlay network 130, such as through the SONET VC provisioning. The mechanism is able to recognize and to process IGMP requests received from the CPEs 104 and 106. The mechanism is further able to perform IP multicast switching based on both the IGMP requests and the multicast channel availability from the SONET egress port, such as the SONET egress ADM 208.

[1032] The point at which IGMP signaling and multicast routing meet is referred to herein as a “rendezvous point”. Making the rendezvous point as close as possible to the CPE is beneficial. For most telcos’ xDSL architectures, the DSLAM 126 is the preferred candidate to perform rendezvous point functions. Alternatively, other devices (e.g. a router) connected to the trunking interface of the DSLAM 126 can perform the rendezvous point functions.

[1033] Because the SONET rings 132 and 134 are provisioned for the delivery of unidirectional broadcast video, and video content is available at an egress ADM interface, the DSLAM 126 or alternative rendezvous point does not need to propagate the IGMP request to other IP multicast routing devices.

[1034] Thus, the act of providing the aforementioned mechanism can comprise adapting a DSLAM to perform the rendezvous point functions.

[1035] As indicated by block 212, broadcast video contents are collected and re-staged in a head-end office of the head-end network 102. The broadcast video contents can be encoded and compressed using any suitable technique, e.g. MPEG2.

[1036] Based on the broadcast video contents, a video payload is encoded and carried by and in an IP multicast envelope, as indicated by block 213. As indicated by block 214, the IP multicast envelope is fed from the head-end network 102 into the SONET ring ingress ADM 202.

[1037] As indicated by block 216, the broadcast video contents in the form of IP multicast envelopes are unidirectionally communicated in the SONET rings 132 and 134. The IP multicast envelopes are communicated from the SONET ring 132 to the SONET ring 134 by the cross connect element 136. The SONET rings 132 and 134 transport the broadcast video contents toward the DSLAM 126 and the CPEs 104 and 106.

[1038] Users of the CPEs 104 and 106 can select any of the broadcast video contents for view. In response to a user selection of a particular video channel, a CPE such as a set-top box generates an IGMP request message for the DSLAM 126. For purposes of illustration and example, consider the CPE 104 generating the IGMP request message.

[1039] As indicated by block 220, the DSLAM 126 receives the IGMP request message from its line interface. In response thereto, the DSLAM 126 looks up or otherwise determines the particular video channel's availability from its trunking interface based on a class-D IP address or a group address provided by the IGMP request message (block 222). If the particular video channel is available, the DSLAM 126 delivers the particular video channel from the trunking interface to its line interface (block 224).

[1040] The particular video channel is communicated from the line interface to the requesting CPE 104. The requesting CPE 104 receives the particular video channel, and decodes and renders video on a display device 230 for view by the user. Examples of the display device 230 include, but are not limited to, a television display and a computer display.

[1041] The aforementioned method ensures an immediate delivery of a broadcast channel and eliminates uncertainties, such as tree-building delays, lack of bandwidth, and unavailability of QoS, in dynamically building a multicast routing tree. Establishing such a tree would involve tree grafting and pruning based on channel request interests.

[1042] The CPEs 104 and 106 may need to send data to the head-end network 102 for interactions or for downloading an Electronic Program Guide (EPG), for example. Such data is well-suited for a unicast network, such as either the legacy data network 100 or the dedicated separate data network 140. FIG. 6 is a flow chart of an embodiment of a method of transporting unicast interactions between the CPEs 104 and 106 and the head-end network 102.

[1043] As indicated by block 240, a CPE such as the CPE 104 issues a unicast request for a destination in the head-end network 102. As indicated by block 242, the DSLAM 126 receives the unicast request via its line interface.

[1044] If the architecture is absent the dedicated separate data network 140 (such as the architecture of FIG. 3), the unicast traffic is directed to the legacy data network 100 (as indicated by block 244). The legacy data network 100 communicates the unicast traffic to the destination in the head-end network 102.

[1045] If the architecture includes the dedicated separate data network 140 (such as the architecture of FIG. 4), the DSLAM 126 decides whether the request should be directed (e.g. routed or switched) to the legacy data network 100 or the dedicated separate data network 140 (as indicated by block 246). The DSLAM 126 makes the decision based on its policy-based routing set-up. Either the legacy data network 100 or the dedicated separate data network 140 communicates the unicast traffic to the head-end network 102.

[1046] As indicated by block 250, the DSLAM 126 directs (e.g. routes or forwards) unicast traffic from the head-end network 102 to its intended destination CPE. The DSLAM 126 can receive the unicast traffic via either the legacy data network 100 or the dedicated separate data network 140. As indicated by block 252, the DSLAM 126 directs

(e.g. routes or forwards) unicast traffic from the Internet 124 to its intended destination CPE without use of the enhanced multicast architecture.

[1047] The head-end network 102 may have its own routing policy to the Internet 124. Therefore, traffic from the head-end network 102 to the Internet 124 should find its own way to the Internet 124.

[1048] In summary, embodiments of the present invention introduce an overlay network dedicated for broadcast video delivery to avoid QoS complexity and to enhance QoS management. The overlay network can comprise a SONET broadcast overlay network provisioned as a unidirectional network. The core of the SONET broadcast overlay network uses a ring topology rather than a tree topology. DSLAMs are adapted to know that the SONET is provisioned as a unidirectional broadcast network. Policy-based mechanisms are built into the DSLAMs to understand that one or more overlay networks exist to serve different transport purposes. Accordingly, each DSLAM makes smart routing/switching decisions based on preset policies including but not limited to broadcast vs. unicast, address range, address group and application types.

[1049] Embodiments of the present invention have the following benefits.

[1050] 1. Since the legacy data network 100 is left largely intact, no expensive network element upgrades are required therein. Since SONET is a cost effective technology, the overlay network 130 is cheaper to implement than expanding the legacy data network 100.

[1051] 2. Neither over-engineering nor ATM/IP QoS interworking or tweaking are required in the legacy data network 100. Accordingly, there are fewer service disruptions.

[1052] 3. The SONET technology is well-understood and widely deployed in and between the telco's central offices (COs). A greater synergy can be created by using an incumbent SONET infrastructure for transporting broadcast video and for other telco applications.

[1053] 4. The SONET static service provisioning used for the overlay network 130 is simpler and more predictable than a dynamic IP-multicast tree pruning and grafting mechanism.

[1054] 5. Since broadcast video applications are asymmetrical and only a unidirectional SONET VC is provisioned, any extra fiber bandwidth can be used either for redundancy protection or for increasing the broadcast service bandwidth.

[1055] 6. Using the overlay network 130 to broadcast video channels does not adversely affect the performance of the legacy data network 100. Traffic fluctuations in the legacy data network 100, does not adversely affect the quality of video channels in the overlay network 130.

[1056] 7. Since the channels are available from the trunking interface of the DSLAM, a channel surfing response time is fast.

[1057] It will be apparent to those skilled in the art that the disclosed embodiments may be modified in numerous ways and may assume many embodiments other than the forms specifically set out and described herein. For example, acts depicted in FIGS. 5 and 6 may be performed in an alternative order. In particular, the acts indicated by blocks 242, 244, 246, 250 and 252 should not be construed as being limited to the order depicted in FIG. 6.

[1058] Accordingly, the above disclosed subject matter is to be considered illustrative, and not restrictive, and the appended claims are intended to cover all such modifications, enhancements, and other embodiments which fall within the true spirit and scope of the present invention. Thus, to the maximum extent allowed by law, the scope of the present invention is to be determined by the broadest permissible interpretation of the following claims and their equivalents, and shall not be restricted or limited by the foregoing detailed description.